REMARKS/ARGUMENTS

The Examiner objects to the Abstract. This objection has been overcome by the rewritten abstract.

The Examiner rejects claims 1 and 2 under 35 U.S.C.§112, second paragraph, as being indefinite. This rejection is now moot in view of the cancellation of claims 1 and 2.

The Examiner rejects claims 1-41 under 35 U.S.C.§102(e) as being anticipated by U.S. Patent 7,200,673 to Augart.

Applicant respectfully traverses the Examiner's rejections.

Augart fails to teach or suggest at least the following italicized limitations of the pending independent claims:

> 12. A method for identifying per-packet load balancing, comprising:
> (a) providing a baseline network topology;
> (b) selecting, from the baseline network topology, first and second addresses associated with first and second routers, respectively, wherein the first router has an associated first hop count relative to a selected node and the second router an associated second hop count relative to the selected node and wherein the first hop count is less than the second hop count;
> (c) transmitting a plurality of test packets from a common source address *to a common selected destination address*, each of the test packets having a time to live equal to or greater than the first hop count;
> (d) receiving a plurality of responses associated with the test packets; and
> (e) *applying the following rules:*
> *(E1) when all of the responses are from a common router, concluding that per-packet load balancing is not in effect; and*
> *(E2) when the responses are from different routers, concluding that per-packet load balancing is in effect.*
>
> 23. A method for identifying per-packet load balancing, comprising:
> (a) providing a baseline network topology;
> (b) selecting, from the baseline network topology, first and second addresses associated with first and second routers, respectively, wherein the first router has an associated first hop count relative to a selected node and the second router an associated second hop count relative to the selected node and wherein the first hop count is less than the second hop count;

(c) transmitting a plurality of test packets from a common source address *to a plurality of differing destination addresses*, each of the test packets having a time to live equal to or greater than the first hop count;
(d) receiving a plurality of responses associated with the test packets; and
(e) *applying the following rules:*

(E1) when all of the responses are from a common router, concluding that at least one of per-destination and per-source/destination load balancing is not in effect; and

(E2) when the responses are from different routers, concluding that at least one of per-destination and per-source/destination load balancing is in effect.

33. A system for detecting load balancing in a distributed processing network, comprising:
(a) a memory comprising a baseline network topology; and
(b) a processor operable to:

(i) select, from the baseline network topology, first and second addresses associated with first and second routers, respectively, wherein the first router has an associated first hop count relative to a selected node and the second router an associated second hop count relative to the selected node and wherein the first hop count is less than the second hop count;

(ii) *transmit first and second sets of test packets, the test packets having a time to live equal to or greater than the first hop count, wherein the first set of test packets are from a common source address to a common selected destination address and the second set of test packets are from a common source address to a plurality of differing destination addresses*;

(iii) receive responses to the first and second sets of test packets; and

(iv) *apply the following rules:*

(A) when all of the responses to the first set of test packets are from a common router, concluding that no per-packet load balancing is in effect;

(B) when the responses to the first set of test packets are from a different routers, concluding that per-packet load balancing is in effect;

(C) when all of the responses to the second set of test packets are from a common router, concluding that at least one of per-destination and per-source/destination load balancing load balancing is not in effect;

(B) when the responses to the second set of test packets are from different routers, concluding that at least one of per-destination and per-source/destination load balancing is in effect.

42. A method, comprising:
(a) providing a set of device addresses associated with a plurality of routers, the plurality of routers being interposed between a testing node and a selected network object;
(b) selecting, from the set of device addresses, a first device address, wherein the first device address is a first hop count from the testing node and a second device address, in the set of device addresses, is a second hop count from the testing node and wherein the first hop count is less than the second hop count;
(c) *transmitting a first set of test packets to at least one of (i) the first device address and (ii) one or more selected destination addresses, each member of the first set of test packets having a Time To Live ("TTL") equal to or greater than the first hop count, wherein the test device on the one hand and the one or more selected destination addresses on the other are located logically on either side of the first device address*;
(d) *transmitting a second set of test packets to multiple destination addresses, each member of the second set of test packets having a TTL equal to or greater than the first hop count, wherein the test device on the one hand and each of the multiple destination addresses on the other are located logically on either side of the first device address*;
(e) receiving a plurality of responses to the first and second sets of test packets;
(f) *applying the following rules:*

(F1) when all of the responses to the first set of test packets are from a router associated with the selected device address, concluding that per-packet load balancing is not in effect;

(F2) when one or more of the responses to the first set of test packets are from a router other than the router associated with the selected device address, concluding that per-packet load balancing is in effect;

(F3) when all of the responses to the second set of test packets are from the router associated with the selected device address, concluding that at least one of per-destination and per-source/destination load balancing is not in effect;

(F4) when one or more of the responses to the second set of test packets are from a router other than the router associated with the selected device address, concluding that at least one of per-destination and per-source/destination load balancing is in effect; and

(g) updating a network topology to reflect the results of steps (e) and (f).

U.S. 7,200,673 to Augart is directed to techniques and apparatuses for determining the *geographic* location of a node on a network. In a representative embodiment, a data packet is received over the network from a second node, the data packet including a network identifier for

the second node and a Time-To-Live (TTL) field that has a value, with the value of the TTL field for the data packet indicating a maximum additional number of hops that could have been made by the data packet. A probe packet addressed to the network identifier for the second node is then sent, the probe packet also including a TTL field. The initial value for the TTL field of the probe packet is set based on the value for the TTL field of the data packet.

By sending a probe packet whose TTL value is based on the TTL value for a received data packet, Augart can identify a router near the originator of the data packet much more quickly than conventional probing techniques would permit. In a preferred embodiment, the number of hops taken by the data packet is estimated based on the TTL field of the data packet. Using this estimated number of hops, one can design a probe packet (e.g., by appropriately setting the initial TTL value of the probe packet) to receive a response from the router immediately prior to the originator of the data packet, or any other desired router along the path. Once a response to the probe packet is received, the response including a network identifier for a router, that network identifier can be compared to a database that includes a geographic location for each of multiple different network identifiers in order to identify a geographic location for the router. If the router is identified to be at a network access point, then in general the requestor can be assumed to be located in the geographic area served by the router. By sending multiple probe packets addressed to the network identifier for the second node, e.g., with initial TTL values for a majority of such probe packets clustered around the estimated number of hops taken by the data packet, Augart is assured of identifying a router that is geographically close to the requester. Moreover, by sending such multiple probe packets without waiting for responses Augart can provide results even faster.

As part of the location process, a check is made for asymmetric routing. At col. 10, lines 12-42, Augart states:

> *In step 153, the host checks to determine whether the responses from the initial set of probe packets indicates asymmetric routing, multi-path routing or any other routing anomaly. Such a situation will occur, for example, if the response corresponding to the probe packet having a TTL value of t did not originate from the requestor or if any response corresponding to a TTL value less than t did originate from the requester.* In an ordinary situation with symmetric

> routing, the response to a probe packet having a TTL value of t (but not the response corresponding to a TTL value of t-1) will be an ICMP Port Unreachable packet. If this is not the case, then it may be determined that either: (i) asymmetric routing, multi-path routing or another routing anomaly has been discovered or (ii) the assumption in step 148 regarding the initial TTL value for the incoming request was wrong. Thus, additional processing may be performed to verify the existence and identify of such anomalous routing. In the event that such a situation is identified, an alternative probing strategy (such as a full "shotgun" approach that runs a modified traceroute program that does not wait between packet transmissions) is performed in step 154 (e.g., using probe packets having TTL values ranging from 1 to 30 or from 1 until an ICMP Port Unreachable packet has been received). Upon completion of such alternative probing operation and verification of an anomalous routing situation or non-standard initial TTL value, the existence of such situation can be stored in a database for future reference (e.g., in step 148 and/or step 150). It is noted that such additional probing for anomaly detection may occur either in real time, offline or using a combination of the two.

(Emphasis supplied.)

As can be seen from this paragraph, Augart does not teach per-packet, per-destination, or per-source/destination load balancing let alone how to distinguish among them. Augart simply determines whether some type of routing anomaly is in effect. A routing anomaly includes not only asymmetric routing but also multi-path routing. To distinguish between two possible reasons for a response not originating from a requestor, namely (i) asymmetric routing, multi-path routing or another routing anomaly has been discovered or (ii) the assumption in step 148 regarding the initial TTL value for the incoming request was wrong, Augart teaches a full "shotgun" approach using differing TTLs in test packets. It says nothing about sending multiple test packets to common or varied destinations. Thus, Augart is unable to determine whether multi-path routing is in effect let alone the type of multi-path routing in effect.

Accordingly, the pending claims are allowable.

The dependent claims provide added reasons for allowance.

By way of example, dependent claims 18, 20, 28, 30, and 40 require the determination of whether, on a selected link, per-packet or at least one of per-destination and per-source/ destination load balancing is in effect.

Dependent claims 19, 29, 41, 47-48, and 55 are directed to identifying asymmetric links. As noted, Augart fails to distinguish between multi-path routing and asymmetric links. Augart further fails to check router tables for asymmetric paths. (S*ee* Specification at pages 17-18.)

Dependent claims 36, 50 and 52-54 are directed to incrementing the TTL for the second router and sending further test packets. Unlike the present invention, Augart does not teach or suggest progressively incrementing the hop count to cause the hop count in the TTL field to be at least equal to the hop count to selected routers. In this way and as currently claimed, the hop count increases the further one moves away from a source node. Link-by-link the present invention can locate and flag load-balanced links. In contrast, Augart specifically limits TTL to a hop count within a geographical proximity to the selected node.

Based upon the foregoing, Applicants believe that all pending claims are in condition for allowance and such disposition is respectfully requested. In the event that a telephone conversation would further prosecution and/or expedite allowance, the Examiner is invited to contact the undersigned.

Respectfully submitted,

SHERIDAN ROSS P.C.

By: Douglas Swartz

Douglas W. Swartz
Registration No. 37,739
1560 Broadway, Suite 1200
Denver, Colorado 80202-5141
(303) 863-9700

Date: Sept 14, 2007

ANNOTATED SHEET

Fig 3A

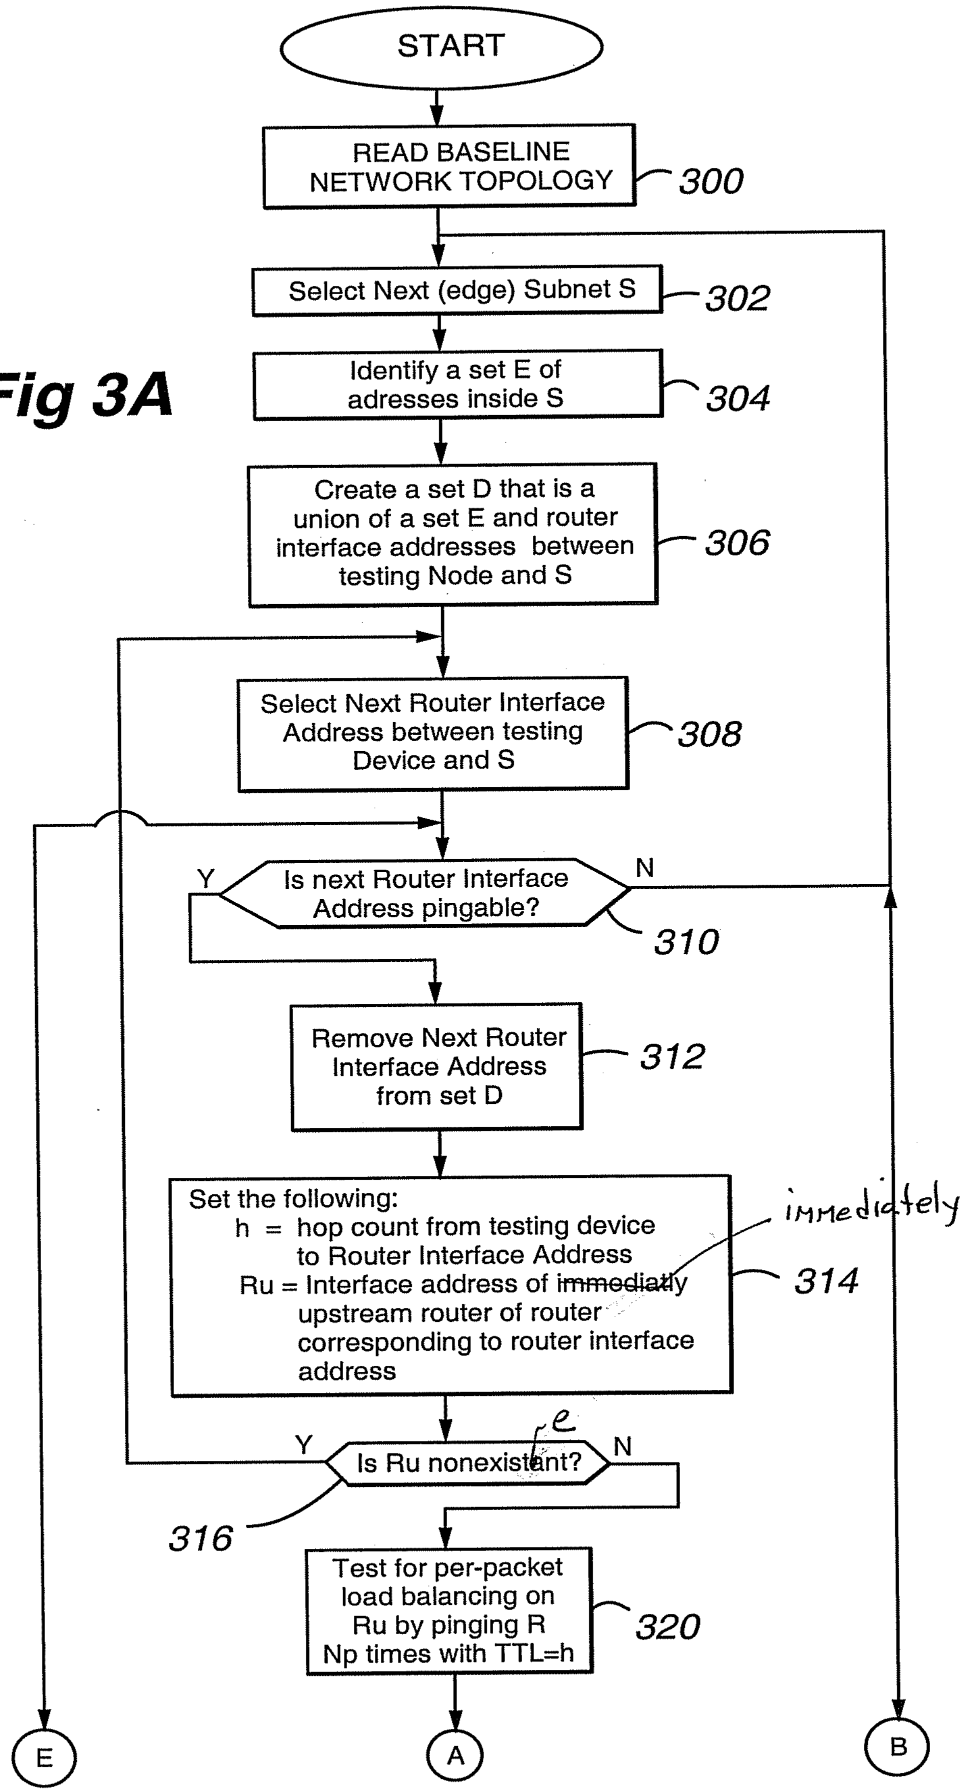

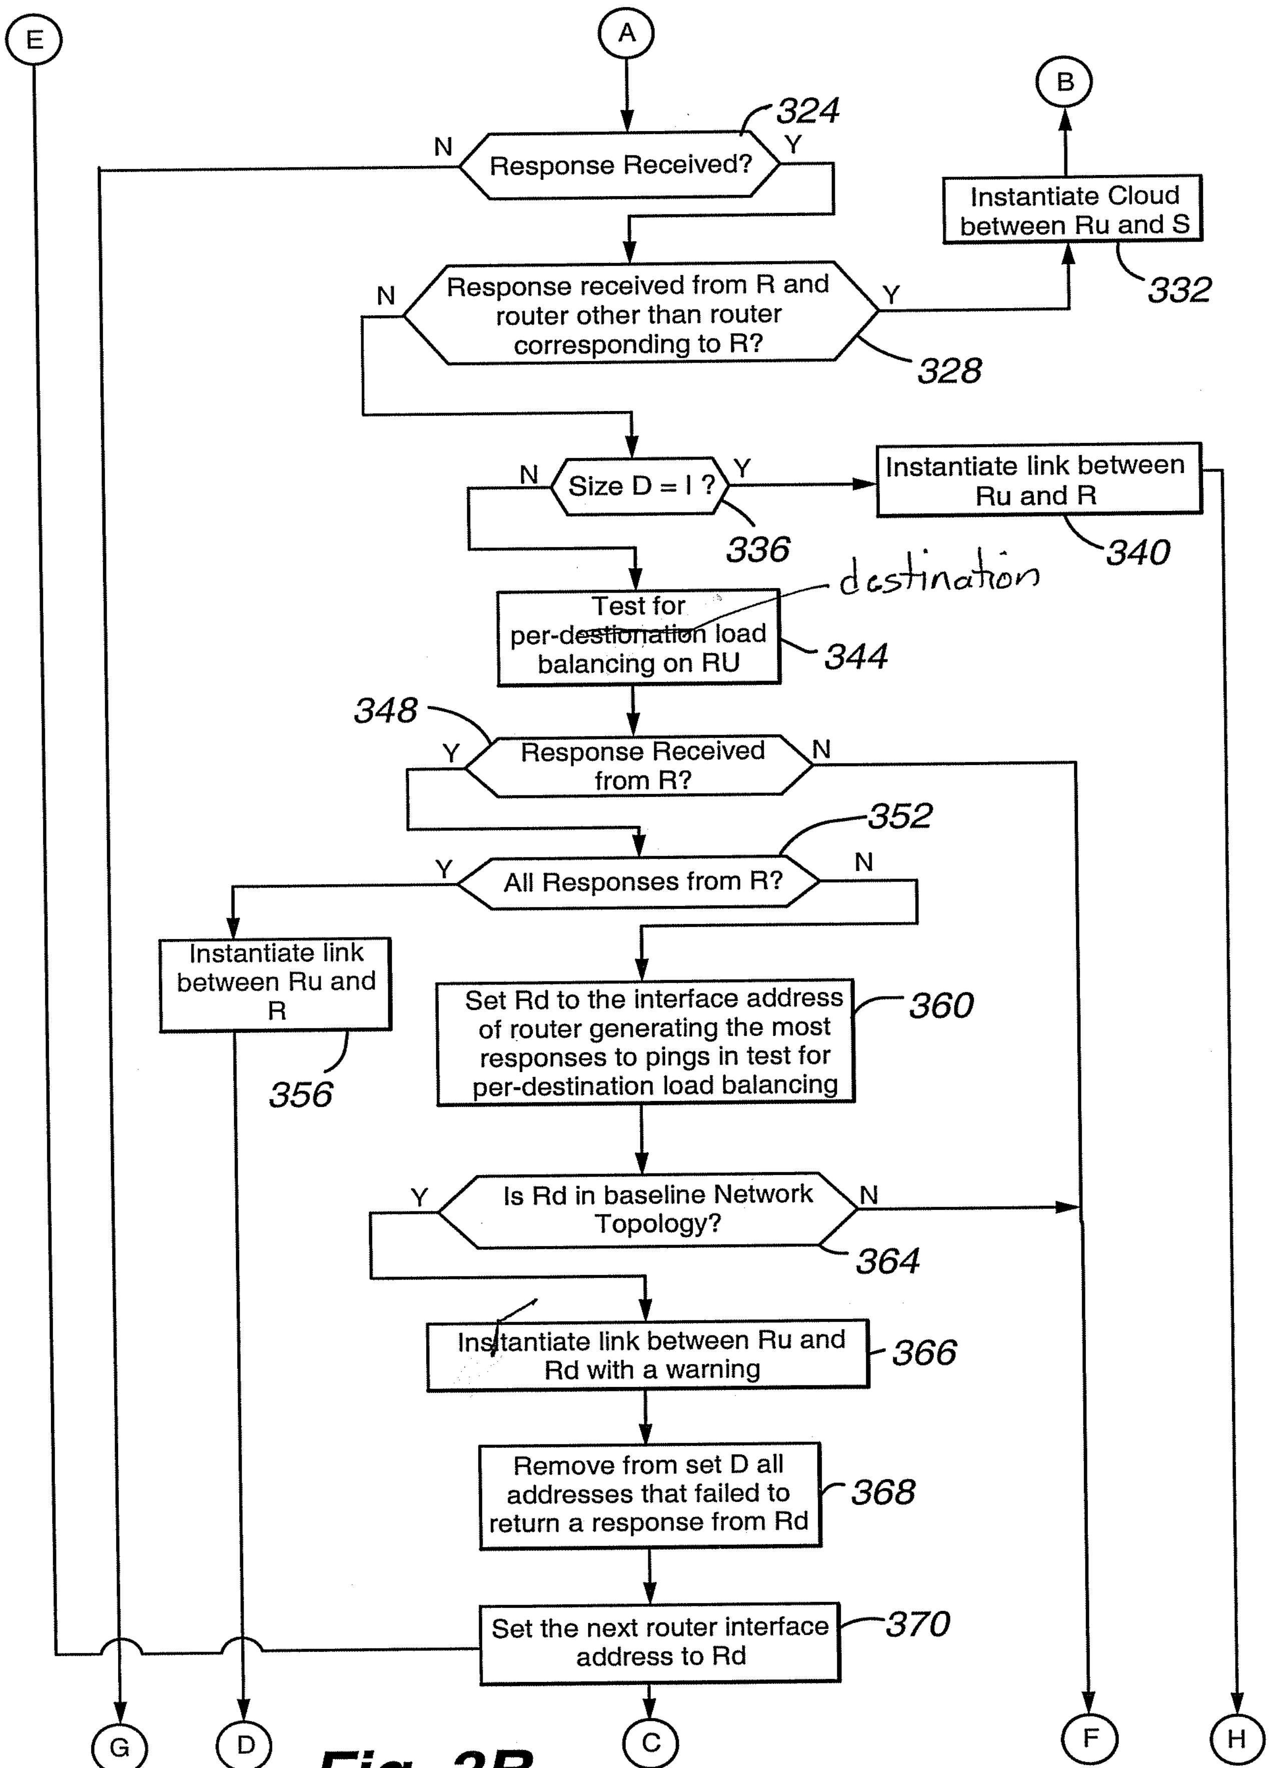

E
A
B
324
N
Response Received?
Y
Instantiate Cloud between Ru and S
332
N
Response received from R and router other than router corresponding to R?
Y
328
N
Size D = I ?
Y
Instantiate link between Ru and R
336
340
destination
Test for per-~~destionation~~ load balancing on RU
344
348
Y
Response Received from R?
N
352
Y
All Responses from R?
N
Instantiate link between Ru and R
360
Set Rd to the interface address of router generating the most responses to pings in test for per-destination load balancing
356
Y
Is Rd in baseline Network Topology?
N
364
Instantiate link between Ru and Rd with a warning
366
Remove from set D all addresses that failed to return a response from Rd
368
Set the next router interface address to Rd
370
G
D
C
F
H

Fig. 3B

ANNOTATED SHEET

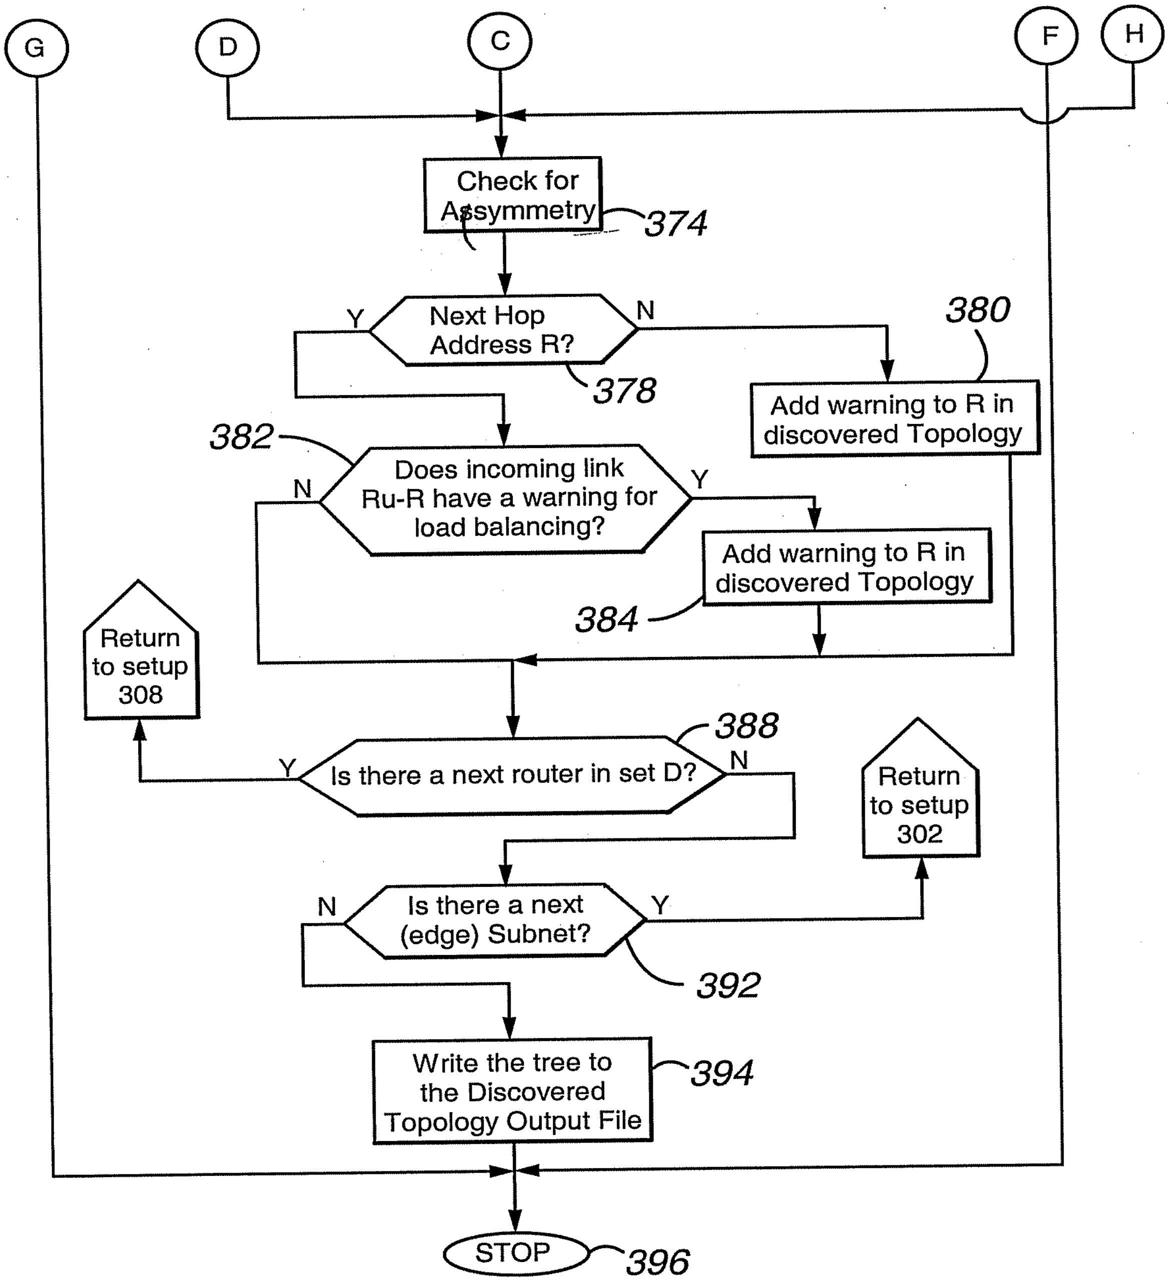

Fig. 3C